



10/ 829677

052

Docket No.: 03226/393001; P8512
(PATENT)

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Letters Patent of:
Matthew A. Ahrens et al.

Patent No.: 7,424,574

Issued: September 9, 2008

For: METHOD AND APPARATUS FOR DYNAMIC
STRIPING

Certificate
NOV 19 2008
of Correction

**REQUEST FOR CERTIFICATE OF CORRECTION
PURSUANT TO 37 CFR 1.322**

Attention: Certificate of Correction Branch
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Dear Sir:

Upon reviewing the above-identified patent, Patentee noted a typographical error which should be corrected.

In the Claims:

In Claim 12, column 10, line 11, the word "string" should be --striping--.

The error was not in the application as filed by applicants; accordingly no fee is required.

Transmitted herewith is a proposed Certificate of Correction effecting such amendment. Also enclosed, as evidence of the error, is a copy of the claims as issued, and a copy of the Claims as allowed. Patentee respectfully solicits the granting of the requested Certificate of Correction.

NOV 19 2008


Patent No.: 7,424,574

Docket No.: 03226/393001; P8512

Applicants believe no fee is due with this request. However, if a fee is due, please charge our Deposit Account No. 50-0591, under Order No. 03226/393001.

Dated: November 12, 2008

Respectfully submitted,

By  #63372 ALY DOSA

Robert P. Lord

Registration No.: 46,479

OSHA · LIANG LLP

909 Fannin Street, Suite 3500

Houston, Texas 77010

(713) 228-8600

(713) 228-8778 (Fax)

NOV 19 2008

scope of the invention as disclosed herein. Accordingly, the scope of the invention should be limited only by the attached claims.

What is claimed is:

1. A method for dynamic striping, comprising:
 - receiving a request to write a file into a storage pool, wherein the file comprises a first data block, a second data block, and a first indirect block;
 - determining a physical disk location in the storage pool to store the first data block using a first dynamic striping policy, wherein the first dynamic striping policy comprises one selected from a group consisting of a dynamic striping policy based on physical disk speed, a dynamic striping policy based on free space available on physical disks, a dynamic striping policy based on load on physical disks, and a round robin policy;
 - storing the first data block using the first dynamic striping policy at the physical disk location in the storage pool;
 - storing the second data block in the storage pool using the first dynamic striping policy;
 - selecting a second dynamic striping policy, wherein the second dynamic striping policy comprises one selected from the group consisting of the dynamic striping policy based on physical disk speed, the dynamic striping policy based on free space available on physical disks, the dynamic striping policy based on load on physical disks, and the round robin policy, and wherein the first dynamic striping policy is distinct from the second dynamic striping policy; and
 - storing the first indirect block using the second dynamic striping policy, wherein the indirect block includes a first block pointer referencing the first data block and a second block pointer referencing the second data block.
2. The method of claim 1, further comprising: retrieving the first data block using the first indirect block.
3. The method of claim 1, further comprising: assembling the first indirect block, wherein assembling the first indirect block comprises populating a block pointer.
4. The method of claim 3, wherein populating the block pointer comprises:
 - storing the first data block checksum in a checksum field within the block pointer; and
 - storing the first data block location in the block pointer, wherein storing the data block location comprises storing a metaslab ID and offset.
5. The method of claim 4, further comprising: storing a birth value in a birth field within the block pointer.
6. The method of claim 3, wherein the first indirect block is assembled using a data management unit.
7. The method of claim 1, wherein the storage pool comprises at least one storage device.
8. The method of claim 1, wherein the storage pool is divided into a plurality of metaslabs.
9. The method of claim 8, wherein each of the plurality of metaslabs is associated with a metaslab ID.
10. The method of claim 9, wherein the first data block location comprises the metaslab ID and an offset.
11. The method of claim 1, wherein storing the first data block comprises using a storage pool allocator.
12. A system for storing a first data block, comprising:
 - a storage pool comprising a file, wherein the file comprises: a first data block and a second data block;
 - a first indirect block; and
 - a parent block referencing the first indirect block, wherein the first indirect block comprises a first block pointer that references the first data block and a second block pointer that references the second data

block, wherein the first block pointer comprises the first data block location and the first data block checksum and the second block pointer comprises the second data block location and the second data block checksum; and

a storage pool allocator configured to:

store the first data block and the second data block in the storage pool using a first dynamic striping policy, wherein the first dynamic striping policy comprises one selected from a group consisting of a dynamic striping policy based on physical disk speed, a dynamic striping policy based on free space available on physical disks, a dynamic striping policy based on load on physical disks, and a round robin policy,

select a second dynamic striping policy after storing the first and second data block, wherein the second dynamic striping policy comprises one selected from the group consisting of the dynamic striping policy based on physical disk speed, the dynamic striping policy based on free space available on physical disks, the dynamic striping policy based on load on physical disks, and the round robin policy, and wherein the first dynamic striping policy is distinct from the second dynamic striping policy, and

store the first indirect block and the parent block in the storage pool using the second dynamic striping policy.

13. The system of claim 12, further comprising: a second indirect block, comprising a first indirect block checksum and a first indirect block location, wherein the storage pool allocator is further configured to store the second indirect block in the storage pool using the second dynamic striping policy.

14. The system of claim 12, further comprising: a data management unit configured to assemble the first indirect block and request the storage pool allocator to store the first indirect block.

15. The system of claim 12, wherein the storage pool comprises at least one storage device.

16. The system of claim 12, wherein the storage pool is divided into a plurality of metaslabs.

17. The system of claim 16, wherein each of the plurality of metaslabs is associated with a metaslab ID.

18. The system of claim 17, wherein the first data block location comprises the metaslab ID and an offset.

19. A computer system for dynamic striping, comprising:

a processor;

a memory;

a storage pool; and

software instructions stored in the memory for enabling the computer system under control of the processor, to:

receive a request to write a file into the storage pool, wherein the file comprises a first data block, a second data block, and a first indirect block;

determine a physical disk location in the storage pool to store the first data block using a first dynamic striping policy, wherein the dynamic striping policy comprises one selected from a group consisting of a dynamic striping policy based on physical disk speed, a dynamic striping policy based on free space available on physical disks, a dynamic striping policy based on load on physical disks, and a round robin policy;

store the first data block using the first dynamic striping policy at the physical disk location in the storage pool;

*

11

store the second data block in the storage pool using the first dynamic striping policy;

select a second dynamic striping policy, wherein the second dynamic striping policy comprises one selected from the group consisting of the dynamic striping policy based on physical disk speed, the dynamic striping policy based on free space available on physical disks, the dynamic striping policy based on load on physical disks, and the round robin policy, and wherein the first dynamic striping policy is distinct from the second dynamic striping policy; and

store the first indirect block using the second dynamic striping policy, wherein the indirect block includes a first block pointer referencing the first data block and a second block pointer referencing the second data block.

20. A network system having a plurality of nodes, comprising:

a storage pool comprising a file, wherein the file comprises: a first data block and a second data block;

a first indirect block; and

a parent block referencing the first indirect block, wherein the first indirect block comprises a first block pointer that references the first data block and a second block pointer that references the second data block, wherein the first block pointer comprises the first data block location and the first data block checksum and the second block pointer comprises the second data block location and the second data block checksum; and

12

a storage pool allocator configured to:

store the first data block and the second data block in the storage pool using a first dynamic striping policy, wherein the first dynamic striping policy comprises one selected from a group consisting of a dynamic striping policy based on physical disk speed, a dynamic striping policy based on free space available on physical disks, a dynamic striping policy based on load on physical disks, and a round robin policy,

store the first indirect block and the parent block using a second dynamic striping policy, wherein the second dynamic striping policy comprises one selected from the group consisting of the dynamic striping policy based on physical disk speed, the dynamic striping policy based on free space available on physical disks, the dynamic striping policy based on load on physical disks, and the round robin policy, and wherein the first dynamic striping policy is distinct from the second dynamic striping policy,

wherein the storage pool is located on any one of the plurality of nodes,

wherein the storage pool allocator is located on any one of the plurality of nodes.

* * * * *

CLAIMS AS ALLOWED

(Re-numbered 1-20)

1. A method for dynamic striping, comprising:
 - receiving a request to write a file into a storage pool, wherein the file comprises a first data block, a second data block, and a first indirect block;
 - determining a physical disk location in the storage pool to store the first data block using a first dynamic striping policy, wherein the first dynamic striping policy comprises one selected from a group consisting of a dynamic striping policy based on physical disk speed, a dynamic striping policy based on free space available on physical disks, a dynamic striping policy based on load on physical disks, and a round robin policy;
 - storing the first data block using the first dynamic striping policy at the physical disk location in the storage pool;
 - storing the second data block in the storage pool using the first dynamic striping policy;
 - selecting a second dynamic striping policy, wherein the second dynamic striping policy comprises one selected from the group consisting of the dynamic striping policy based on physical disk speed, the dynamic striping policy based on free space available on physical disks, the dynamic striping policy based on load on physical disks, and the round robin policy, and wherein the first dynamic striping policy is distinct from the second dynamic striping policy; and
 - storing the first indirect block using the second dynamic striping policy, wherein the indirect block includes a first block pointer referencing the first data block and a second block pointer referencing the second data block.
2. The method of claim 1, further comprising:
 - retrieving the first data block using the first indirect block.
3. The method of claim 1, further comprising:
 - assembling the first indirect block, wherein assembling the first indirect block comprises populating a block pointer.

4. The method of claim 3, wherein populating the block pointer comprises:
storing the first data block checksum in a checksum field within the block pointer; and
storing the first data block location in the block pointer, wherein storing the data block location comprises storing a metaslab ID and offset.
5. The method of claim 4, further comprising:
storing a birth value in a birth field within the block pointer.
6. The method of claim 3, wherein the first indirect block is assembled using a data management unit.
7. The method of claim 1, wherein the storage pool comprises at least one storage device.
8. The method of claim 1, wherein the storage pool is divided into a plurality of metaslabs.
9. The method of claim 8, wherein each of the plurality of metaslabs is associated with a metaslab ID.
10. The method of claim 9, wherein the first data block location comprises the metaslab ID and an offset.
11. The method of claim 1, wherein storing the first data block comprises using a storage pool allocator.
12. A system for storing a first data block, comprising:
a storage pool comprising a file, wherein the file comprises:
a first data block and a second data block;
a first indirect block; and
a parent block referencing the first indirect block, wherein the first indirect block comprises a first block pointer that references the first data block and a second block pointer that references the second data block, wherein the first block pointer comprises the first data block location and the first data block

checksum and the second block pointer comprises the second data block location and the second data block checksum; and
a storage pool allocator configured to:

store the first data block and the second data block in the storage pool using a first dynamic striping policy, wherein the first dynamic striping policy comprises one selected from a group consisting of a dynamic striping policy based on physical disk speed, a dynamic striping policy based on free space available on physical disks, a dynamic striping policy based on load on physical disks, and a round robin policy,

select a second dynamic striping policy after storing the first and second data block, wherein the second dynamic striping policy comprises one selected from the group consisting of the dynamic striping policy based on physical disk speed, the dynamic striping policy based on free space available on physical disks, the dynamic striping policy based on load on physical disks, and the round robin policy, and wherein the first dynamic striping policy is distinct from the second dynamic striping policy, and

store the first indirect block and the parent block in the storage pool using the second dynamic striping policy.

13. The system of claim 12, further comprising:

a second indirect block, comprising a first indirect block checksum and a first indirect block location,

wherein the storage pool allocator is further configured to store the second indirect block in the storage pool using the second dynamic striping policy.

14. The system of claim 12, further comprising:

a data management unit configured to assemble the first indirect block and request the storage pool allocator to store the first indirect block.

15. The system of claim 12, wherein the storage pool comprises at least one storage device.
16. The system of claim 12, wherein the storage pool is divided into a plurality of metaslabs.
17. The system of claim 16, wherein each of the plurality of metaslabs is associated with a metaslab ID.
18. The system of claim 17, wherein the first data block location comprises the metaslab ID and an offset.
19. A computer system for dynamic striping, comprising:
 - a processor;
 - a memory;
 - a storage pool; and
 - software instructions stored in the memory for enabling the computer system under control of the processor, to:
 - receive a request to write a file into the storage pool, wherein the file comprises a first data block, a second data block, and a first indirect block;
 - determine a physical disk location in the storage pool to store the first data block using a first dynamic striping policy, wherein the dynamic striping policy comprises one selected from a group consisting of a dynamic striping policy based on physical disk speed, a dynamic striping policy based on free space available on physical disks, a dynamic striping policy based on load on physical disks, and a round robin policy;
 - store the first data block using the first dynamic striping policy at the physical disk location in the storage pool;
 - store the second data block in the storage pool using the first dynamic striping policy;
 - select a second dynamic striping policy, wherein the second dynamic striping policy comprises one selected from the group consisting of the dynamic striping policy based on physical disk speed, the dynamic striping policy based on free space available on physical disks, the dynamic striping policy based on load on physical

disks, and the round robin policy, and wherein the first dynamic striping policy is distinct from the second dynamic striping policy; and
store the first indirect block using the second dynamic striping policy, wherein the indirect block includes a first block pointer referencing the first data block and a second block pointer referencing the second data block.

20. A network system having a plurality of nodes, comprising:

a storage pool comprising a file, wherein the file comprises:

a first data block and a second data block;

a first indirect block; and

a parent block referencing the first indirect block, wherein the first indirect block comprises a first block pointer that references the first data block and a second block pointer that references the second data block, wherein the first block pointer comprises the first data block location and the first data block checksum and the second block pointer comprises the second data block location and the second data block checksum; and

a storage pool allocator configured to:

store the first data block and the second data block in the storage pool using a first dynamic striping policy, wherein the first dynamic striping policy comprises one selected from a group consisting of a dynamic striping policy based on physical disk speed, a dynamic striping policy based on free space available on physical disks, a dynamic striping policy based on load on physical disks, and a round robin policy,

store the first indirect block and the parent block using a second dynamic striping policy, wherein the second dynamic striping policy comprises one selected from the group consisting of the dynamic striping policy based on physical disk speed, the dynamic striping policy based on free space available on physical disks, the dynamic striping policy based on load on physical disks, and the round robin policy, and wherein the first dynamic striping policy is distinct from the second

dynamic striping policy,
wherein the storage pool is located on any one of the plurality of nodes,
wherein the storage pool allocator is located on any one of the plurality of nodes.

**UNITED STATES PATENT AND TRADEMARK OFFICE
CERTIFICATE OF CORRECTION**Page 1 of 1

PATENT NO. : 7,424,574
APPLICATION NO. : 10/828,677
ISSUE DATE : September 9, 2008
INVENTOR(S) : Matthew A. Ahrens et al.

It is certified that an error appears or errors appear in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

In the Claims:


In Claim 12, column 10, line 11, the word "string" should be --striping--.

MAILING ADDRESS OF SENDER (Please do not use customer number below):

Aly Z. Dossa
OSHA · LIANG LLP
909 Fannin Street, Suite 3500
Houston, Texas 77010

1

NOV 19 2008



PTO/SB/92 (10-08)
Approved for use through 11/30/2008. OMB 0651-0031
Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

Application No. (if known): 10/828,677

Attorney Docket No.: 03226/393001; P8512

Certificate of Mailing under 37 CFR 1.8

I hereby certify that this correspondence is being deposited with the United States Postal Service with sufficient postage as first class mail in an envelope addressed to:

Attention: Certificate of Correction Branch
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

on November 12, 2008
Date

Signature

Blanca E. Ramos

Typed or printed name of person signing Certificate

Registration Number, if applicable

(713) 228-8600

Telephone Number

Note: Each paper must have its own certificate of mailing, or this certificate must identify each submitted paper.

Request for Certificate of Correction (No Fee) with attachments (10 pages)
Certificate of Correction (1 page)
Return Receipt Post Card

NOV 19 2008